Journal of Machine Learning Research ? (????) ??-?? Submitted ?/??; Published ?/??

Joint Causal Inference on Observational and Experimental Datasets

Sara Magliacane SARA.MAGLIACANE@GMAIL.COM
VU Amsterdam
Amsterdam, The Netherlands

Institute for Informatics, University of Amsterdam
Amsterdam, The Netherlands

Tom Claassen TOMC@CS.RU.NL
Radboud University Nijmegen
Nijmegen, The Netherlands

Institute for Informatics, University of Amsterdam
Amsterdam, The Netherlands

Joris M. Mooij J.M.MOOIJ@UVA.NL
Institute for Informatics, University of Amsterdam
Amsterdam, The Netherlands

Editor: ??

arXiv:1611.10351v1 [cs.LG] 30 Nov 2016

Abstract

We introduce Joint Causal Inference (JCI), a powerful formulation of causal discovery over multiple datasets that allows to jointly learn both the causal structure and targets of interventions from statistical independences in pooled data. Compared with existing constraint-based approaches for causal discovery from multiple data sets, JCI offers several advantages: it allows for several different types of interventions, it can learn intervention targets, it systematically pools data across different datasets which improves the statistical power of independence tests, and it improves on the accuracy and identifiability of the predicted causal relations. A technical complication that arises in JCI are the occurrence of faithfulness violations due to deterministic relations. We propose a simple but effective strategy for dealing with this type of faithfulness violations. We implement it in ACID, a determinism-tolerant extension of Ancestral Causal Inference (ACI) (Magliacane et al., 2016), a recently proposed logic-based causal discovery method that improves reliability of the output by exploiting redundant information in the data. We illustrate the benefits of JCI with ACID with an evaluation on a simulated dataset.

1. Introduction

Discovering causal relations from data is at the foundation of the scientific method (Ducheyne, 2006). Traditionally, causal relations are either recovered from experimental data in which the variable of interest is perturbed, or from observational data, e.g., using the seminal PC and FCI algorithms (Spirtes et al., 2000; Zhang, 2008).

In the recent literature there have been several proposals for combining observational and experimental data to discover causal relations, showing that this combination can improve greatly on the accuracy and identifiability of the predicted causal relations. Some of the proposed methods are *score-based* (e.g., Eaton and Murphy, 2007; Hauser and Bühlmann, 2012), i.e., they evaluate

©???? Sara Magliacane, Tom Claassen and Joris M. Mooij.

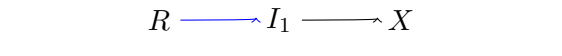

$$X \not\perp_d I_1 \,|\, R \qquad I_1 = f(R) \implies X \perp\!\!\!\perp I_1 \,|\, R$$

Figure 1: **Example of faithfulness violations due to deterministic relations**: Let $I_1 = f(R)$ for a function f. In the graph X cannot be d-separated from I_1 by R. On the other hand, $X \perp\!\!\!\perp I_1 \,|\, R$, because I_1 is fully determined by R, as follows from the properties of conditional independence.

models using a penalized likelihood score, while others (e.g., Hyttinen et al., 2014; Triantafillou and Tsamardinos, 2015; Peters et al., 2015; Magliacane et al., 2016) are *constraint-based*, i.e., they use statistical independences to express constraints over possible models.

In this work we propose Joint Causal Inference (JCI), a formulation of causal discovery over multiple datasets in which we jointly learn both the causal structure and targets of interventions from independence test results. A similar approach was already proposed for score-based methods by Eaton and Murphy (2007), but here we show how it can be extended such that constraint-based methods can be employed. Our goal is to unify the idea of joint inference from observational and experimental data from (Eaton and Murphy, 2007) with the advantages that constraint-based methods have over score-based methods, namely, the ability to handle latent confounders naturally, and, especially in the case of logic-based methods, an easy integration of background knowledge.

Existing constraint-based methods for multiple datasets typically learn the causal structure on each dataset separately and then merge the learnt structures. Typically, the merging process depends on the type of interventions, and most existing methods support only perfect interventions on known targets, e.g., (Hyttinen et al., 2014; Triantafillou and Tsamardinos, 2015). Instead, JCI: (1) allows for several different types of interventions and learns intervention targets; (2) systematically pools data across different datasets, which improves the statistical power of independence tests; and (3) improves the accuracy and identifiability of the predicted causal relations.

On the other hand, JCI poses challenges for current constraint-based methods because of their susceptibility to violations of the *Causal Faithfulness assumption*. Specifically, JCI typically leads to faithfulness violations due to deterministic relations between intervention variables, which cannot be handled by standard constraint-based methods. A simple example of this type of violation is shown in the graph in Figure 1. We propose a simple but effective strategy for dealing with this type of faithfulness violations. We implement it in ACID (Ancestral Causal Inference with Determinism), a determinism-tolerant extension of ACI (Ancestral Causal Inference) (Magliacane et al., 2016), a recently proposed logic-based causal discovery method that improves reliability of the output by exploiting redundant information in the data. In our evaluation on synthetic data we show that JCI with ACID improves on the accuracy of the causal predictions with respect to simply merging separately learned causal graphs.

2. Preliminaries

In this Section we review a few useful concepts from the related work and introduce the notation we use in the rest of the paper. In the following, we represent variables with uppercase letters, while sets of variables are denoted by boldface.

2.1 Graph terminology

Throughout the paper we assume that the data generating process can be modeled by a causal Directed Acyclic Graph (DAG) that may contain latent variables. For simplicity, we do not consider selection bias. A directed edge $X \to Y$ in the causal DAG represents a direct causal relationship of cause X on effect Y. We say that X is a *parent* of Y, and denote the set of parents of Y as $\text{PA}(Y)$. A sequence of directed edges $X_1 \to X_2 \to \cdots \to X_n$ is a *directed path*. If there is a directed path from X to Y (or $X = Y$) then X is an *ancestor* of Y (denoted as $X \dashrightarrow Y$). We denote the set of ancestors of Y as $\text{AN}(Y)$. If there is no directed path from X to Y (and $X \neq Y$) we denote this as $X \not\dashrightarrow Y$.

For a set of variables $\boldsymbol{W}$, we extend the definition of parents $\text{PA}(\boldsymbol{W})$ to the union of all parents of any variable $W \in \boldsymbol{W}$. We define the set of ancestors of $\boldsymbol{W}$, $\text{AN}(\boldsymbol{W})$, similarly. In other words, we write $X \dashrightarrow \boldsymbol{W}$ if there exists at least one effect $Y \in \boldsymbol{W}$ that has X as an ancestor, i.e., $X \dashrightarrow Y$. Similarly, we write $X \not\dashrightarrow \boldsymbol{W}$ if $X \not\dashrightarrow Y$ for all $Y \in \boldsymbol{W}$.

2.2 Independences and d-separation

For disjoint sets of random variables $\boldsymbol{X}, \boldsymbol{Y}, \boldsymbol{W}$, distributed according to a joint probability distribution P, we denote the conditional independence of $\boldsymbol{X}$ and $\boldsymbol{Y}$ given $\boldsymbol{W}$ in P as $\boldsymbol{X} \perp\!\!\!\perp \boldsymbol{Y} \mid \boldsymbol{W}\ [P]$, and conditional dependence as $\boldsymbol{X} \not\perp\!\!\!\perp \boldsymbol{Y} \mid \boldsymbol{W}\ [P]$. We often omit $[P]$ when it is obvious from the context which probability distribution we are referring to. We call the cardinality $|\boldsymbol{W}|$ the *order* of the conditional (in)dependence relation.

A well-known graphical criterion for DAGs with many implications for causal discovery is *d-separation* (Pearl, 2009; Spirtes et al., 2000):

Definition 1 *For disjoint sets of variables $\boldsymbol{X}, \boldsymbol{Y}, \boldsymbol{W}$ in a DAG $\mathcal{G}$, we say that $\boldsymbol{X}$* is d-separated from $\boldsymbol{Y}$ by $\boldsymbol{W}$ in $\mathcal{G}$, *written $X \perp_d Y \mid \boldsymbol{W}\ [\mathcal{G}]$, iff for every path π in $\mathcal{G}$ that connects any $X \in \boldsymbol{X}$ with any $Y \in \boldsymbol{Y}$, at least one of the following holds:*

- *π contains a collider not in $\text{AN}(\boldsymbol{W})$, or*
- *π contains a non-collider in $\boldsymbol{W}$.*

The opposite, i.e., d-connection, is denoted as $\boldsymbol{X} \not\perp_d \boldsymbol{Y} \mid \boldsymbol{W}\ [\mathcal{G}]$. We often omit $[\mathcal{G}]$ from the notation when it is obvious which DAG are we referring to.

There are three assumptions that have been thoroughly discussed in literature (see for example (Spirtes et al., 2000; Peters et al., 2014)) connecting conditional independences in the observational distribution P and d-separation in the underlying causal DAG $\mathcal{G}$. For disjoint (sets) of variables $X, \boldsymbol{X}, \boldsymbol{Y}, \boldsymbol{W}$:

- *Causal Markov Assumption*: d-separation in the causal DAG $\mathcal{G}$, implies conditional independence in the observational distribution P. We can also rewrite this implication using the negated concepts as conditional dependence implies d-connection, i.e., $\boldsymbol{X} \not\perp\!\!\!\perp \boldsymbol{Y} \mid \boldsymbol{W}\ [P] \implies \boldsymbol{X} \not\perp_d \boldsymbol{Y} \mid \boldsymbol{W}\ [\mathcal{G}]$
- *Causal Faithfulness Assumption*: $\boldsymbol{X} \perp\!\!\!\perp \boldsymbol{Y} \mid \boldsymbol{W}\ [P] \implies \boldsymbol{X} \perp_d \boldsymbol{Y} \mid \boldsymbol{W}\ [\mathcal{G}]$,
- *Causal Minimality Assumption*: $\forall Q \in \text{PA}(X) : X \not\perp\!\!\!\perp Q \mid \text{PA}(X) \setminus \{Q\}$ (Peters et al., 2014).

$Y = f(X) \implies Z \perp\!\!\!\perp W \,|\, X$ — Functionally determined example

$Y = X \oplus U \implies Z \perp\!\!\!\perp W \,|\, \{U, Y\}$ — Non-functionally determined example

Figure 2: **Faithfulness violations examples**: In the left example, $Z \perp\!\!\!\perp W \,|\, X$, because $Z \perp\!\!\!\perp W \,|\, Y$ and Y is determined by X, even though $Z \not\perp\!\!\!\perp_d W \,|\, X$. In the right example, let $Y = X \oplus U$, where X and U are binary variables and $\oplus$ is the XOR function. Conditioning on $\{U, Y\}$ fully determines X, even if $\{U, Y\}$ are not ancestors of X. This means that $Z \perp\!\!\!\perp W \,|\, \{U, Y\}$, although $Z \not\perp\!\!\!\perp_d W \,|\, \{U, Y\}$.

If we assume both the Causal Markov and Causal Faithfulness assumptions to hold, the conditional independences of the observational distribution correspond one-to-one with the d-separations in the causal DAG. This setting is very favourable for causal discovery, thus both of these assumption are usually made in most constraint-based approaches.

2.3 Deterministic relations and faithfulness violations

Although often reasonable, the Causal Faithfulness assumption is violated in some cases, notably in the common case of deterministic relations among variables, e.g., for Structural Causal Model equations (Pearl, 2009) in which there is no noise term. We show two examples in Figure 2, where we represent equations without noise terms with blue edges.

Some of the faithfulness violations related to determinism are captured by an extension to the d-separation criterion, the *D-separation* criterion, first presented in (Geiger et al., 1990) and later extended in (Spirtes et al., 2000). Under the Causal Markov and Minimality assumptions, the formulation of D-separation presented in (Geiger et al., 1990) is proven to be complete for the restricted setting where determinism arises only due to *functionally determined* relations, defined recursively as variables that are fully determined by their parents, or more precisely:

Definition 2 *A variable X is* functionally determined *by a set $\boldsymbol{W}$ (for a given graph $\mathcal{G}$) if $X \in \boldsymbol{W}$, or X is a deterministic function of its parents* $\text{PA}(X)$ *in $\mathcal{G}$ and* $\text{PA}(X) \subseteq \boldsymbol{W}$.

In Section 3.8 of their book, Spirtes et al. (2000) extend D-separation to model also some deterministic relations that are due to variables being determined when conditioning on their non-ancestors. In the right example in Figure 2 we show an example in which the definition of D-separation from (Geiger et al., 1990) fails to capture the faithfulness violation that is due to the deterministic relation between non-ancestors of X.

Although the version of D-separation by Spirtes et al. (2000) retains completeness for the restricted case of functionally determined relations and under the Minimality assumption, it is not proven to be complete in general. Nevertheless, Spirtes et al. (2000) introduce several useful concepts for handling general deterministic relations, so we summarize their findings here, adapting them to our notation. We start with the assumption that we have complete knowledge of all deterministic relations in the system.

Assumption 1 *$\mathcal{D}$ is a complete list of all the deterministic relations among variables, where each entry in the list $\langle\{V_1, \ldots, V_{n-1}\}, V_n\rangle$ indicates that variable V_n is a deterministic function of variables $\{V_1, \ldots, V_{n-1}\}$, but it is not of any strict subset of $\{V_1, \ldots, V_{n-1}\}$.*

Note that the deterministic relations in the system are not restricted only to the "functionally determined" relations as in Definition 2 (i.e., Structural Causal Model equations with no noise terms). For example, in Figure 2 left $\mathcal{D} = \langle\{X\}, Y\rangle$, while in Figure 2 right, $\mathcal{D} = \langle\{X, U\}, Y\rangle$, $\langle\{Y, X\}, U\rangle$, $\langle\{Y, U\}, X\rangle$. We use $\mathcal{D}$ to define a function that returns all determined variables for a given set of variables $\boldsymbol{W}$.

Definition 3 *Given a set of variables $\boldsymbol{W}$ and a complete list of deterministic relations $\mathcal{D}$, we define* $\text{DET}(\boldsymbol{W}, \mathcal{D})$ *as the set of variables determined according to $\mathcal{D}$ by (a subset of) $\boldsymbol{W}$. We omit $\mathcal{D}$, using only* $\text{DET}(\boldsymbol{W})$*, if it is obvious from the context which list are we referring to.*

Note that $\text{DET}(\boldsymbol{W}, \mathcal{D})$ trivially includes $\boldsymbol{W}$ itself. Also, any variable with constant value is by definition in $\text{DET}(\boldsymbol{W}, \mathcal{D})$ for all $\boldsymbol{W}$ as it is determined by $\emptyset$. Following Spirtes et al. (2000), we can use $\text{DET}()$ to extend d-separation for deterministic relations.

Definition 4 *Given a DAG $\mathcal{G}$, three disjoint sets of variables $\boldsymbol{X}, \boldsymbol{Y}, \boldsymbol{W}$, and the complete list of deterministic relations $\mathcal{D}$ in the distribution P, we define $\boldsymbol{X}$ and $\boldsymbol{Y}$ as* D-separated *by $\boldsymbol{W}$ w.r.t. $\mathcal{D}$ and $\mathcal{G}$, denoted as $\boldsymbol{X} \perp_D \boldsymbol{Y} \mid \boldsymbol{W}$ $[\mathcal{D}, \mathcal{G}]$, iff for every path π in $\mathcal{G}$ between any $X \in \boldsymbol{X}$ and any $Y \in \boldsymbol{Y}$, at least one of the following holds:*

- *π contains a collider not in* $\text{AN}(\boldsymbol{W})$*, or*
- *π contains a non-collider that is in* $\text{DET}(\boldsymbol{W}, \mathcal{D})$*, or*
- *the first or last node of π is in* $\text{DET}(\boldsymbol{W}, \mathcal{D})$.

If $\text{DET}(\boldsymbol{W}, \mathcal{D}) = \boldsymbol{W}$*, D-separation reduces to standard d-separation. We omit $\mathcal{D}$ and $\mathcal{G}$ if it is obvious from the context which list and graph are we referring to.*

Under the Causal Markov Assumption, this formulation of D-separation is proven to imply independence (Spirtes et al., 2000), or more precisely: $\boldsymbol{X} \perp_D \boldsymbol{Y} \mid \boldsymbol{W} \implies \boldsymbol{X} \perp\!\!\!\perp \boldsymbol{Y} \mid \boldsymbol{W}$. For the case of functionally determined relations and under the Minimality assumption, this version of D-separation is complete, i.e., it correctly describes all additional independences due to functionally determined relations, because it reduces to the version of D-separation by Geiger et al. (1990).

3. Related work

Given a set of observational and interventional datasets, most constraint-based methods learn the causal structure from each dataset separately and then merge the learned structures (Hyttinen et al., 2014; Triantafillou and Tsamardinos, 2015). This conservative viewpoint derives from a misplaced caution in applying conditional independence tests on mixtures of distributions, as would happen if one were to naively pool the data, as discussed by Eberhardt (2008).

Merging causal structures learnt on each dataset separately has several drawbacks with respect to a method that can jointly use all datasets, as for example the score-based method from Eaton and Murphy (2007). First, the merging approaches require known targets and cannot learn the targets

of the interventions, since this type of information is only available when using multiple datasets jointly. Moreover, they cannot take advantage of certain interventional datasets, e.g., in the case of a single data point per interventional setting, as for example happens in a popular genomics dataset (Kemmeren et al., 2014). Two other drawbacks are less statistical power because of the separation into smaller isolated datasets, and, as we will show with some examples in Section 4.1, less identifiable relations with respect to a joint method.

There is some preliminary work on some special cases in which to apply constraint-based methods with mixtures of observational and experimental datasets, e.g., (Eberhardt, 2008; Lagani et al., 2012; Peters et al., 2015; Borboudakis and Tsamardinos, 2016), but the problem hasn't been yet systematically discussed and formalized in a general framework. In particular, no work addresses learning the intervention targets, possibly jointly with the causal graph, from independence constraints.

Some approaches describe sufficient, although quite restrictive, conditions under which pooling data does not change the conditional distribution of the variables under consideration. In particular, Eberhardt (2008) describes how naively pooling data from different experimental settings, while discarding the information of which experimental setting a sample was taken, may give wrong results. Thus he proposes a sufficient condition that allows one to pool data for a given independence test when the conditional distribution of the tested variables is the same in all experimental conditions. Lagani et al. (2012) present two other approaches: (i) perform the conditional independence tests separately in each dataset, then define the pooled dependence as a disjunction of the single dependences; (ii) pool experimental conditions that differ only in the value of at most one intervened variable. Both of these approaches describe restrictive conditions in which one can pool datasets, while in this paper we argue that, when done systematically, e.g., as we will show in the next Section, one can in principle always pool all available datasets.

Other approaches like (Peters et al., 2015) focus on certain specific combinations of independence tests that are performed jointly on all datasets. Invariant Causal Prediction (ICP) (Peters et al., 2015) is a causal discovery method that can exploit different experimental settings, returning a *conservative subset* of ancestors (or parents, if one assumes there are no latent confounders) for a given target variable Y. The main assumption is that the conditional distribution of Y given its parents does not change in the different interventional settings (also referred to as invariance or modularity in (Pearl, 2009) and (Spirtes et al., 2000)), unless Y is directly intervened upon. Since the method searches for predictors that are invariant across different settings, it can safely pool together a subset of settings in a new virtual "experimental" setting to increase the statistical power for settings with few data. On the other hand, as we show with some examples in Section 4.2, the conservativeness of the ICP estimates sometimes reduces significantly the causal information we can derive.

Borboudakis and Tsamardinos (2016) recently proposed ETIO, a constraint-based approach that can learn the structure over multiple datasets and supports perfect ("structural") and soft ("parametric") interventions on known ("not uncertain") targets. Soft interventions are represented with an indicator variable for each target variable, and are treated as standard variables by the causal discovery method. The idea of this work is closely related to our paper, but it does not consider several important details. For example, in the case of multiple targets, there are several intervention variables that may be deterministically related to each other. In general, also other deterministic relations can be present, leading to faithfulness violations that are not taken into account in ETIO. Instead, in this paper we propose a detailed formalization of this type of issues, including a strategy to deal with certain types of faithfulness violations that arise in this context. Moreover, we allow

R	I_1	I_2	X_1	X_2	X_4
1	20	0	0.1	0.2	0.5
1	20	0	0.13	0.21	0.49
1	20	0	...	...	...
2	20	1	...	...	...
3	30	0	...	...	...
4	30	1	...	...	...
5	30	2	...	...	...

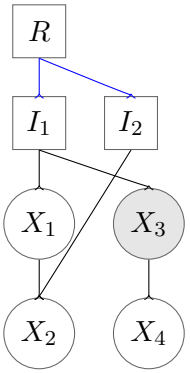

Figure 3: A set of five experimental datasets in a raw data form (left) and as a causal DAG (right) representing the causal structure of the system variables $X_1, \ldots, X_4$, the regime variable R and the intervention variables I_1, I_2. The intervention variable I_1 could represent for example the temperature at which each experiment was performed, while the intervention variable I_2 may represent the dosage of a drug added in some of the experiments. The left part of the matrix is the experimental design matrix. Similarly to the causal influence diagrams from (Dawid, 2002), we represent system variables $\{X_j\}_{j\in\mathcal{X}}$ as circles, which are filled for hidden variables, while we represent the regime R and the intervention variables I_1, I_2 as squares.

each intervention variable to be attached to multiple (possibly unknown) targets, which enables the joint learning of the intervention targets and causal structure.

4. Joint Causal Inference (JCI)

We propose to model jointly in a single causal graph several observational or experimental datasets $\{D_r\}_{r\in\{1\ldots n\}}$, each representing the data after a (possibly empty) set of interventions on one or more, possibly unknown, targets. We assume that there is a unique underlying causal DAG in all of these datasets, defined over the same set of variables that we call the *system variables*, $\{X_j\}_{j\in\mathcal{X}}$, some of which are possibly hidden. This assumption precludes certain types of interventions, notably, perfect interventions (Pearl, 2009). On the other hand, it allows for many other types of interventions, e.g., soft interventions (Markowetz et al., 2005), mechanism changes (Tian and Pearl, 2001), fat-hand interventions (Eaton and Murphy, 2007), activity interventions (Mooij and Heskes, 2013), etc., as long as they do not induce new (in)dependences, which can be seen as modifications to the underlying DAG.

Each dataset D_r has an associated joint probability distribution $P_r((X_j)_{j\in\mathcal{X}})$. Using the terminology from (Dawid, 2002), we call the different distributions in the datasets *regimes*. In related work different names have been used, e.g. *experimental conditions* or *environments* (Mooij and Heskes, 2013; Peters et al., 2015). We introduce two types of dummy variables in the data:

- a **regime variable** R, representing which dataset D_r a data point is from, i.e., $\forall r = 1 \ldots n$, $R = r$ for data from D_r.

- **intervention variables** $\{I_i\}_{i\in\mathcal{I}}$, which are deterministic functions of the regime R. Intuitively, intervention variables represent the interventions performed in each dataset. In ab-

sence of any information on the interventions performed in the datasets, we can use as intervention variables the indicator variables for each of the datasets.

We describe an example in Figure 3.

We assume that the causal relations between system variables $\{X_j\}_{j\in\mathcal{X}}$ and the introduced dummy variables R and $\{I_i\}_{i\in\mathcal{I}}$ can be represented as an acyclic Structural Causal Model (SCM) with jointly independent exogenous variables $\{E_k\}_{k\in\mathcal{X}\cup\{R\}}$:

$$\begin{aligned}
&\begin{cases} R &= E_R, \\ I_i &= g_i(R), \quad i \in \mathcal{I}, \\ X_j &= f_j(X_{\text{PA}_{\mathcal{X}}(X_j)}, I_{\text{PA}_{\mathcal{I}}(X_j)}, E_j), \quad j \in \mathcal{X}, \end{cases} \\
&P\big((E_k)_{k\in\mathcal{X}\cup\{R\}}\big) = \prod_{k\in\mathcal{X}\cup\{R\}} P(E_k).
\end{aligned} \tag{1}$$

Here, $\text{PA}_{\mathcal{X}}(X_j)$ are the system variable parents of X_j, while $\text{PA}_{\mathcal{I}}(X_j)$ denote its intervention parents and E_j is the exogenous parent of X_j.

We represent this SCM with a causal DAG $\mathcal{G}$ such that the Causal Markov and Minimality assumptions hold in $\mathcal{G}$. We show an example in Figure 3, where we model five datasets with the same underlying causal structure. For clarity, we graphically distinguish between regime/intervention variables (using rectangles) and system variables (using circles), similarly to the causal influence diagrams from (Dawid, 2002).

Our modeling assumptions are applicable when the value of the regime/intervention variables are determined by the experimenter *before* the system variables are measured. In other words, *the system variables cannot cause the regime/intervention variables*. In addition, we assume that the values set by the experimenter are chosen independently of any other possible cause of the system variables. In other words, we assume there to be no latent confounders between the regime variable and the system variables. For the purposes of causal discovery, there is nothing else that really distinguishes the regime/intervention variables from the system variables: they can both be considered to be random variables, where the distribution of the regime/intervention variables just reflects the empirical distribution of the choices that the experimenter has made.

Intervention variables are functions of the regime variable, and do not have any associated noise. This means that they are *determined* by the regime. We represent these functions as a matrix:

Definition 5 *We define the **experimental design matrix** as the matrix representing the functional relations between R and each intervention variable I_i.*

In general, other deterministic relations between dummy variables may arise. For example, consider the example in Table 1 left in which $I_{Akt-Inh}$ that represents an inhibitor that was added when the regime is an odd number, while I_{U0126} indicates another inhibitor that was added when the regime is an even number. These two variables determine each other. Even though this is clear from the experimental design matrix, it is not visible in the causal influence diagram.

In this paper, we focus on a special case, assuming that the only deterministic relations are the regime R determining each of the intervention variables $\{I_i\}_{i\in\mathcal{I}}$. For this restricted case, all deterministic relations are *functionally determined* relations, and therefore D-separation is sound and complete under the Causal Markov and Minimality assumption (Geiger et al., 1990). We show an example of the experimental design matrix for the functionally determined case in Table 1 right,

R	$I_{Akt-Inh}$	I_{U0126}
1	0	1
2	1	0
3	0	1
4	1	0

R	$I_{Akt-Inh(\neg U0126)}$	I_1
1	0	0
2	1	0
3	0	1
4	1	1

Table 1: Example of experimental design matrix with additional deterministic relations beyond the functionally determined ones (left) and a reduced version with only functionally determined relations (right). In the right version we join $I_{Akt-Inh}$ and I_{U0126} in a single intervention variable representing the addition of a single inhibitor at a time. We also add the intervention variable I_1 so we can distinguish between datasets $R = 1$ and $R = 3$ (and similarly for $R = \{2, 4\}$).

which represents the "functionalization" of the left example matrix. We conjecture that for the more general case of deterministic relations between dummy variables, e.g. Table 1 left, D-separation as defined in (Spirtes et al., 2000) is complete, and thus JCI can be still applied, but we leave the proof for future work.

We are interested in the completeness of D-separation, because we wish to use it to relax the standard Causal Faithfulness assumption. In our setting this assumption would be too restrictive, so we relax it to allow for violations due to deterministic relations between the regime and the intervention variables. We define our relaxed version, that we call *D-Faithfulness assumption*, as follows: for three disjoint sets of variables $\boldsymbol{X}, \boldsymbol{Y}, \boldsymbol{W}$ and a probability distribution P that satisfies both the Causal Markov assumption for $\mathcal{G}$ and the list of deterministic relations $\mathcal{D}$, we assume that $\boldsymbol{X} \perp\!\!\!\perp \boldsymbol{Y} \mid \boldsymbol{W}\ [P] \implies \boldsymbol{X} \perp_D \boldsymbol{Y} \mid \boldsymbol{W}\ [\mathcal{D}, \mathcal{G}]$, where $\perp_D$ represents D-separation as defined in Definition 4. This assumption, in conjunction with the previous ones, implies that in JCI independences correspond one-to-one with D-separations.

Given all of the mentioned assumptions, we define *Joint Causal Inference (JCI)* as the discovery of the causal DAG $\mathcal{G}$ that represents jointly all datasets and intervention variables from independence test results. Moreover, we call any causal discovery method that can solve a JCI instance a *JCI method*. We will show in Section 4.2 that the ideas behind some previous approaches, e.g., (Peters et al., 2015), can be seen as special cases of JCI.

4.1 Joint Causal Inference improves on the identifiability w.r.t. merging learnt structures

As already mentioned, formulating causal inference on multiple datasets as JCI offers several advantages with respect to the approaches in which the causal graphs are learnt separately from each dataset and then merged, e.g. (Triantafillou and Tsamardinos, 2015; Hyttinen et al., 2014). One of the advantages is the improved identifiability. In this Section, we show a few examples, where, for simplicity, we assume oracle inputs and model only the regime variable.

In Figure 4 we show a simple example in which JCI improves identifiability. In the absence of information on the intervention targets, we cannot identify the causal direction between the variables when we learn the structure separately and combine it. In the same case, a JCI method is able to correctly reconstruct the causal structure by using additional conditional independence tests with the regime variable, specifically, $R \not\perp\!\!\!\perp X$, $R \not\perp\!\!\!\perp Y$, $R \perp\!\!\!\perp Y \mid X$, $R \not\perp\!\!\!\perp X \mid Y$ and $X \not\perp\!\!\!\perp Y \mid R$. Using the background knowledge from JCI that system variables cannot cause the regime R, i.e.,

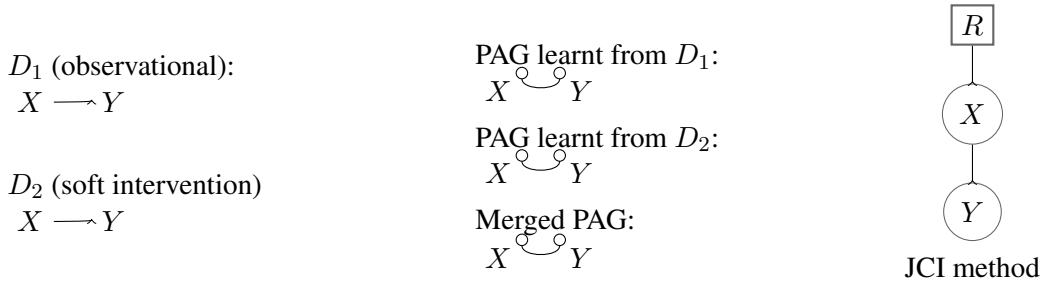

Figure 4: **A simple example in which JCI improves identifiability:** Consider two datasets with the same underlying DAG, one of which has a soft intervention (left example). If we learn the causal graphs of X and Y in each dataset separately and then merge them, e.g. as described in (Triantafillou and Tsamardinos, 2015), we cannot learn the causal direction but only that they are dependent (middle example). JCI adds more variables and thus conditional independence tests that allow to distinguish the direction (right example, see details in the text).

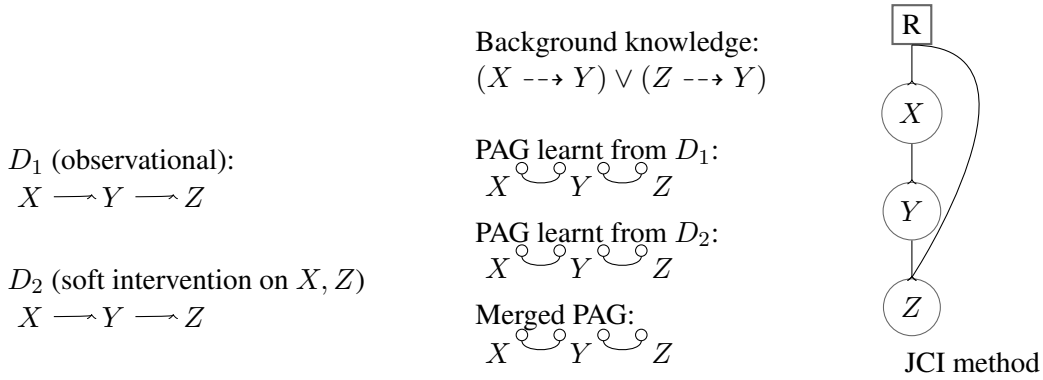

Figure 5: **A more complex example in which JCI improves identifiability:** If we have background knowledge on the intervention targets, e.g. we know that in one of the datasets X and Z are intervened upon, we can use this information to extract some extra background knowledge in the form of ancestral relations. In general, this is not enough information to decide the causal structure. Instead, a JCI method can identify the causal structure (right example, see details in the text).

$X \not\dashrightarrow R \wedge Y \not\dashrightarrow R$, and that there are no latent confounders between R and the system variables, we can infer $R \to X$ and $X \to Y$ with any JCI method supporting direct causal relations.

If the targets of the intervention are known for each dataset it is possible to retrieve their descendants (and non-descendants) by checking which variables change in each dataset with respect to the observational case. In ACI (Magliacane et al., 2016) this technique was successfully applied by performing a two-sample Kolmogorov-Smirnov test on the changed variables and using its confidence as a weight for the resulting weighted ancestral statement. For example, if one were to know that X is the intervention target in D_2 in the example in Figure 4, the change in Y in D_2 with respect to D_1 would imply that $X \dashrightarrow Y$.

Although (weighted) ancestral relations help in simple cases as the previous example, in general they cannot reproduce the same results as JCI, e.g. in the example in Figure 5. In this case, knowing that in dataset D_2 the intervention targets are X and Z, and observing that Y changes significantly,

Figure 6: **Two examples in which ICP is too conservative in the JCI setting:** In the left example, for variable Y ICP finds several sets $\boldsymbol{S}$ that satisfy $R \perp\!\!\!\perp Y \mid \boldsymbol{S}$, e.g. $\{X\}$, $\{Z\}$ and $\{X, Z\}$, but their intersection is $\emptyset$. Instead, a JCI method that supports direct causal relations will correctly infer that the single parent of Y is Z. In the right example, naively adding intervention variables does not allow to estimate the ancestors of a variable Y, which would otherwise be estimated correctly.

allows us to reconstruct that one of these targets causes Y, which is not enough to reconstruct the causal graph. Instead, a JCI method that supports direct causal relations can take advantage of the additional conditional independence tests with the regime variable. Similarly to the previous example, a JCI method can infer $R \to X$ and $X \to Y$, and, with the same reasoning, $R \not\perp\!\!\!\perp Z$ implies that $R \to Z$. Moreover, from $X \not\perp\!\!\!\perp Z \mid R$, $X \perp\!\!\!\perp Z \mid R \cup Y$ we can infer that $Y \dashrightarrow Z$, and, since there is no set of variables $\boldsymbol{S}$ such that $Y \perp\!\!\!\perp Z \mid \boldsymbol{S}$, this causal relation is direct, i.e., $Y \to Z$.

4.2 Reformulation of related work as special cases of JCI

Local Causal Discovery (LCD) (Cooper, 1997) is a simple algorithm on an observational dataset that searches for variables X, Y, W that satisfy the pattern $W \to X \to Y$, where W is an uncaused variable. A simple trick to apply LCD to multiple observational and experimental datasets with soft interventions is to use R as W, since the regime variable is by definition uncaused. Then LCD can be summarized as:

$$(R \not\perp\!\!\!\perp X) \wedge (X \not\perp\!\!\!\perp Y) \wedge (R \perp\!\!\!\perp Y \mid X) \implies (X \dashrightarrow Y) \tag{2}$$

This rule can be seen as a restricted case of the JCI setting, in which we can iteratively pick pairs of variables (X, Y) and apply the above rule to detect a subset of the causal graph.

An extension of this approach is used in Invariant Causal Prediction (ICP) (Peters et al., 2015). ICP can also model the regime variable R (which is called the *discrete environment variable* in that work), but it does not model the intervention variables. An example is shown in the Appendix of (Peters et al., 2015). Given a target variable Y that is not directly intervened upon, we can reformulate the main idea behind ICP as the search for the intersection of all the sets $\boldsymbol{S}$ such that $R \perp\!\!\!\perp Y \mid \boldsymbol{S}$:

$$\boldsymbol{S}^* = \bigcap_{\boldsymbol{S}: R \perp\!\!\!\perp Y \mid \boldsymbol{S}} \boldsymbol{S}. \tag{3}$$

In the absence of confounders, $\boldsymbol{S}^*$ is a conservative estimate of a subset of the parents of Y, even when the Causal Faithfulness assumption is violated. If we cannot exclude the presence of confounders, as in the JCI setting, then ICP requires the Causal Faithfulness assumption to provide $\boldsymbol{S}^*$ as an estimate of a subset of the ancestors of Y. We can see that this reformulation of ICP as a special case of JCI that extends LCD with a more conservative estimate. In principle, one could easily integrate the conservative estimate (3) in a JCI method to provide more accurate estimates for the top predictions, but we leave this for future work.

On the other hand, depending on the set of interventions in the available datasets, ICP may be overly conservative compared to JCI. Besides the restriction on the variable Y not to be directly intervened upon in any dataset, which is not necessary in JCI, there are some other cases in which ICP provides an overly conservative estimate of the set of ancestors. We show two examples in Figure 6. Specifically, in the left example the estimated set of ancestors for a variable Y that is two hops away from the intervened variable X is empty, while a JCI method that supports direct causal relations can find the correct parent set Z. Similarly, as shown in the right example, naively adding the intervention variables reduces the applicability of ICP to only estimate ancestors of variables that are directly intervened upon (e.g., Z). This naive addition would allow ICP to learn the intervention targets, but not the structure of the causal graph.

5. A strategy for extending constraint-based methods for JCI

Joint Causal Inference provides some challenges for current constraint-based methods:

- faithfulness violations due to deterministic relations between the dummy variables,
- the availability of complex background knowledge on the dummy variables that can improve structure learning and recover from some of the faithfulness violations, e.g. R can only cause a system variable through an intervention variable.

There is some work on dealing with faithfulness violations in the PC algorithm (Lemeire et al., 2012), but it assumes causal sufficiency (in our context, no hidden variables in $\mathcal{G}$), and cannot handle complex background knowledge. Other constraint-based algorithms, specifically, logic-based causal discovery methods, e.g., (Hyttinen et al., 2014; Triantafillou and Tsamardinos, 2015; Magliacane et al., 2016) can handle complex background knowledge and causal insufficiency, but cannot deal with faithfulness violations due to deterministic relations.

Here we propose a simple but effective strategy for dealing with faithfulness violations due to deterministic relations. We rephrase the constraints of a constraint-based algorithm in terms of d-separations and d-connections, instead of independence test results. At testing time we decide for each independence test result which d-separations and d-connections can be soundly derived from it and provide them as input to the modified constraint-based algorithm.

Before introducing the rules that we use to derive sound d-separation and d-connections from input independence test results, we first summarise the basic properties of conditional independence originally introduced in (Dawid, 1979), which we will use to prove an intermediate Lemma. We follow the notation and ordering from a more recent publication (Constantinou and Dawid, 2015):

Proposition 6 *Let X, Y, Z, W be random variables. We write $W \preceq Y$ to denote that W is a function of Y, or in other words $W = f(Y)$ for a measurable function f. Then the following properties hold:*

1. $X \perp\!\!\!\perp Y \mid Z \iff Y \perp\!\!\!\perp X \mid Z$,
2. $X \perp\!\!\!\perp Y \mid X$,
3. $X \perp\!\!\!\perp Y \mid Z$ *and* $W \preceq Y \implies X \perp\!\!\!\perp W \mid Z$,
4. $X \perp\!\!\!\perp Y \mid Z$ *and* $W \preceq Y \implies X \perp\!\!\!\perp Y \mid (Z, W)$,

5. $X \perp\!\!\!\perp Y \,|\, Z \wedge X \perp\!\!\!\perp W \,|\, (Y, Z) \implies X \perp\!\!\!\perp (Y, W) \,|\, Z$.

Lemma 7 *For $X, Y, \boldsymbol{W}$ disjoint (sets) of random variables and random variable $F \preceq \boldsymbol{W}$:*

$$X \perp\!\!\!\perp Y \,|\, \boldsymbol{W} \iff X \perp\!\!\!\perp Y \,|\, (\boldsymbol{W}, F). \tag{4}$$

Proof This is a simple consequence of the properties of conditional independence that we reviewed in Proposition 6. We first show one direction of the implication:

$$X \perp\!\!\!\perp (\boldsymbol{W}, Y) \,|\, (\boldsymbol{W}, Y) \text{ (Property 2)} \implies X \perp\!\!\!\perp F \,|\, (\boldsymbol{W}, Y) \text{ (Property 3)}$$

$$\begin{aligned} X \perp\!\!\!\perp Y \,|\, \boldsymbol{W} \wedge X \perp\!\!\!\perp F \,|\, (\boldsymbol{W}, Y) &\implies X \perp\!\!\!\perp (Y, F) \,|\, \boldsymbol{W} \text{ (Property 5)} \\ &\implies X \perp\!\!\!\perp (Y, F) \,|\, (\boldsymbol{W}, F) \text{ (Property 4)} \\ &\implies X \perp\!\!\!\perp Y \,|\, (\boldsymbol{W}, F) \text{ (Property 3)} \end{aligned}$$

Then we show also the other direction:

$$\begin{aligned} X \perp\!\!\!\perp Y \,|\, (\boldsymbol{W}, F) \wedge X \perp\!\!\!\perp F \,|\, \boldsymbol{W} &\implies X \perp\!\!\!\perp (F, Y) \,|\, \boldsymbol{W} \text{ (Property 5)} \\ &\implies X \perp\!\!\!\perp Y \,|\, \boldsymbol{W} \text{ (Property 3)} \end{aligned}$$

■

Our proposed strategy for dealing with faithfulness violations due to deterministic relations is based the following assumptions:

Assumption 2 *We assume that we have a complete list of deterministic relations $\mathcal{D}$ and that the deterministic relations in $\mathcal{D}$ are restricted to* functionally determined *relations. We also assume that the Causal Markov, Minimality and D-Faithfulness assumption hold.*

These assumptions are a subset of those used in JCI. Similarly to JCI, the restriction to functionally determined relations is used to provide a justification for the D-Faithfulness assumption based on the completeness of D-separation in that setting. As for JCI, we conjecture that D-separation may be complete in a more general setting, making this strategy applicable also more generally, but we leave the proof for future work.

Given these assumptions, we can now use Lemma 7 to prove a sound conversion between D-separation statements to d-separation statements:

Theorem 8 *Let $X, Y, \boldsymbol{W}$ be disjoint (sets) of random variables and $\text{DET}(\boldsymbol{W})$ the variables determined by (a subset of) $\boldsymbol{W}$ (see Section 2). Under Assumption 2 the following holds:*

$$\text{If } X, Y \notin \text{DET}(\boldsymbol{W}): \quad X \perp_D Y \,|\, \boldsymbol{W} \iff X \perp_d Y \,|\, \text{DET}(\boldsymbol{W}), \tag{5}$$

where $\perp_d$ is standard d-separation *and $\perp_D$ is* D-separation.[1]

1. Note that if X or $Y \in \text{DET}(\boldsymbol{W})$, then d-separation is not defined.

Proof The core of the proof is based on Lemma 7 and Assumption 2:

$$X \perp_D Y \,|\, \boldsymbol{W} \iff X \perp\!\!\!\perp Y \,|\, \boldsymbol{W} \iff X \perp\!\!\!\perp Y \,|\, \textsc{Det}(\boldsymbol{W}) \text{ by Lemma (7)}$$
$$\iff X \perp_D Y \,|\, \textsc{Det}(\boldsymbol{W}) \iff X \perp_d Y \,|\, \textsc{Det}(\boldsymbol{W}).$$

The first equivalence follows from the Causal Markov and D-Faithfulness assumptions, while the second is based on Lemma (7). The third equivalence also follows from the assumptions, while the last one is based on the definition of D-separation, which reduces to d-separation when conditioning on a set $\boldsymbol{W} = \textsc{Det}(\boldsymbol{W})$. ■

Using the result from Theorem 8, we can now introduce our strategy for dealing with faithfulness violations due to deterministic relations. First we rephrase a constraint-based algorithm in terms of d-separations and d-connections, which is usually a trivial change, as shown in Section 6. Then we can convert the problem of possibly unfaithful independences to the problem of possibly incomplete input. Specifically, we can derive a subset of sound d-separations and d-connections from independence test results as follows:

Corollary 9 *Let $X, Y, \boldsymbol{W}$ be disjoint (sets) of random variables and $\textsc{Det}(\boldsymbol{W})$ the variables determined by (a subset of) $\boldsymbol{W}$ (see Section 2). Under Assumption 2 the following holds:*

- $X \not\perp\!\!\!\perp Y|\boldsymbol{W} \implies X \not\perp_d Y|\boldsymbol{W}$,
- *If* $X, Y \notin \textsc{Det}(\boldsymbol{W})$: $\quad X \perp\!\!\!\perp Y \,|\, \boldsymbol{W} \implies X \perp_d Y \,|\, \textsc{Det}(\boldsymbol{W})$,

where $\perp_d$ is standard d-separation*, $\not\perp_d$ is d-connection.*

Proof The first implication follows from the Causal Markov assumption, while the second follows from the Causal Markov, Minimality and D-Faithfulness assumptions, and Theorem 8. ■

Note that this procedure outputs d-separations only for a subset of independence test results, ignoring independences when X or $Y \in \textsc{Det}(\boldsymbol{W})$.

The simple strategy in Corollary 9 can be applied to any constraint-based method, providing that it can deal with partial inputs, i.e. missing results for a certain independence test. Logic-based methods as (Hyttinen et al., 2014; Magliacane et al., 2016) can be run out-of-the-box with partial inputs, while other standard algorithms like FCI (Zhang, 2008) would require possibly complicated extensions. Anytime FCI (Colombo et al., 2012) allows one to ignore (in)dependences above a certain order, but up to that order they are all required to be available.

6. Ancestral Causal Inference With Determinism (ACID)

We implement the strategy in Corollary 9 in Ancestral Causal Inference with Determinism (ACID) as a determinism-tolerant extension of Ancestral Causal Inference (ACI), a recently introduced logic-based method described in (Magliacane et al., 2016).

6.1 ACI

ACI accurately reconstructs ancestral structures ("indirect" causal relations), also in the presence of latent variables and statistical errors. Ancestral structures are formally defined as:

Definition 10 *An **ancestral structure** is any relation on the observed variables that satisfies the non-strict partial order axioms:*

$$(\text{reflexivity}) : X \dashrightarrow X \tag{6}$$

$$(\text{transitivity}) : X \dashrightarrow Y \ \wedge \ Y \dashrightarrow Z \implies X \dashrightarrow Z \tag{7}$$

$$(\text{antisymmetry}) : X \dashrightarrow Y \wedge Y \dashrightarrow X \implies X = Y \tag{8}$$

The underlying causal DAG induces a unique ancestral structure on the observed variables: the transitive closure of the direct causal relations (directed edges) in the DAG.

ACI encodes the ancestral structure definition and six other simple rules:

Lemma 11 *For X, Y, Z, U, $\boldsymbol{W}$ disjoint (sets of) variables:*

1. $(X \perp\!\!\!\perp Y \mid \boldsymbol{W}) \wedge (X \not\dashrightarrow \boldsymbol{W}) \implies X \not\dashrightarrow Y$,
2. $(X \perp\!\!\!\perp Y \mid \boldsymbol{W}) \wedge (X \not\perp\!\!\!\perp Y \mid \boldsymbol{W} \cup Z) \implies (X \not\perp\!\!\!\perp Z \mid \boldsymbol{W}) \wedge (Z \not\dashrightarrow \{X, Y\} \cup \boldsymbol{W})$,
3. $(X \not\perp\!\!\!\perp Y \mid \boldsymbol{W}) \wedge (X \perp\!\!\!\perp Y \mid \boldsymbol{W} \cup Z) \implies (X \not\perp\!\!\!\perp Z \mid \boldsymbol{W}) \wedge (Z \dashrightarrow \{X, Y\} \cup \boldsymbol{W})$,
4. $X \not\perp\!\!\!\perp Y \mid \boldsymbol{W}) \wedge (X \perp\!\!\!\perp Y \mid \boldsymbol{W} \cup Z) \wedge (X \perp\!\!\!\perp Z \mid \boldsymbol{W} \cup U) \implies X \perp\!\!\!\perp Y \mid \boldsymbol{W} \cup U$,
5. $(Z \not\perp\!\!\!\perp X \mid \boldsymbol{W}) \wedge (Z \not\perp\!\!\!\perp Y \mid \boldsymbol{W}) \wedge (X \perp\!\!\!\perp Y \mid \boldsymbol{W}) \implies X \not\perp\!\!\!\perp Y \mid \boldsymbol{W} \cup Z$.

These rules are shown to be sound assuming the Causal Markov and Causal Faithfulness assumptions. Causal discovery is then reformulated as an optimization problem where a loss function is optimized over possible ancestral structures. Given a list of weighted inputs, e.g. a set of conditional independences weighted by their confidence, the loss function sums the weights of all the inputs that are violated in a candidate ancestral structure. In addition, ACI provides a method for scoring causal predictions, which roughly approximates their marginal probability.

6.2 ACID

Out-of-the-box ACI is not able to deal with the faithfulness violations due to deterministic relations, and thus cannot be used for JCI. Therefore, we propose *Ancestral Causal Inference with Determinism (ACID)*, which extends ACI following the strategy discussed in Section 5. We reformulate the logical rules of ACI in terms of *d-separation*, completely decoupling them from any assumption on the relation between (in)dependences and d-separations/connections, e.g., Causal Faithfulness. These new rules, that we call the *ACID rules*, are almost identical to the original ACI rules, except that the independences $\perp\!\!\!\perp$ are substituted by d-separations $\perp_d$, and the dependences $\not\perp\!\!\!\perp$ by d-connections $\not\perp_d$, as we show in the following:

Lemma 12 *For X, Y, Z, U, $\boldsymbol{W}$ disjoint (sets of) variables:*

1. $(X \perp_d Y \mid \boldsymbol{W}) \wedge (X \not\dashrightarrow \boldsymbol{W}) \implies X \not\dashrightarrow Y$
2. $(X \perp_d Y \mid \boldsymbol{W}) \wedge (X \not\perp_d Y \mid \boldsymbol{W} \cup Z) \implies (X \not\perp_d Z \mid \boldsymbol{W}) \wedge (Z \not\dashrightarrow \{X, Y\} \cup \boldsymbol{W})$
3. $(X \not\perp_d Y \mid \boldsymbol{W}) \wedge (X \perp_d Y \mid \boldsymbol{W} \cup Z) \implies (X \not\perp_d Z \mid \boldsymbol{W}) \wedge (Z \dashrightarrow \{X, Y\} \cup \boldsymbol{W})$
4. $(X \not\perp_d Y \mid \boldsymbol{W}) \wedge (X \perp_d Y \mid \boldsymbol{W} \cup Z) \wedge (X \perp_d Z \mid \boldsymbol{W} \cup U) \implies X \perp_d Y \mid \boldsymbol{W} \cup U$

5. $(Z \not\perp_d X \mid \boldsymbol{W}) \wedge (Z \not\perp_d Y \mid \boldsymbol{W}) \wedge (X \perp_d Y \mid \boldsymbol{W}) \implies X \not\perp_d Y \mid \boldsymbol{W} \cup Z$

Proof The proofs of these rules are slight modifications of the proofs of the ACI rules. For completeness, we provide them in the Appendix. ■

Then at testing time, we provide as inputs to ACID only the sound d-separations and d-connections that can be derived with Corollary 9. Similarly to other logic-based methods, the ACID rules are sound also with partial inputs (i.e. when any d-separation information may not be available). On the other hand, using partial inputs may reduce the completeness of causal discovery. We consider this a minor issue, since our focus is on prediction accuracy, and ACI is already known not to be complete in the general case, but nevertheless can obtain state-of-the-art accuracies (Magliacane et al., 2016).

6.3 ACID-JCI

To improve the identifiability and accuracy of the predictions, we also add as background knowledge a series of logical rules describing the causal structure of the regime and intervention variables that apply in the JCI setting:

Lemma 13 *For any set of variables $\boldsymbol{W}$:*

1. $\forall j, k \in \mathcal{X}$ *s.t.* $j \neq k$ *and* $X_j, X_k \notin \boldsymbol{W}: \; X_j \perp_d X_k \,|\, \boldsymbol{W} \implies X_j \perp_d X_k \,|\, \boldsymbol{W} \cup R$*, i.e., adding R to the separating set cannot open paths,*

2. $\forall i \in \mathcal{I}, \forall j, k \in \mathcal{X}$ *s.t.* $j \neq k$ *and* $X_j, X_k \notin \boldsymbol{W}: \; X_j \perp_d X_k \,|\, \boldsymbol{W} \implies X_j \perp_d X_k \,|\, \boldsymbol{W} \cup I_i$*, i.e., adding an intervention variable to the separating set cannot open paths,*

3. $\forall i \in \mathcal{I}$ *s.t.* $I_i \notin \boldsymbol{W}: \; (R \dashrightarrow I_i) \wedge (R \not\perp_d I_i \,|\, \boldsymbol{W})$*, i.e., R is a direct cause of each I_i,*

4. $\forall i \in \mathcal{I}, j \in \mathcal{X}: \; (X_j \not\dashrightarrow R) \wedge (X_j \not\dashrightarrow I_i)$*, i.e., system variables cannot cause any I_i or R,*

5. $\forall j \in \mathcal{X}$ *s.t.* $X_j \notin \boldsymbol{W}: R \dashrightarrow X_j \implies (\exists i \in \mathcal{I}$ *s.t.* $I_i \dashrightarrow X_j) \wedge (R \perp_d X_j \,|\, \boldsymbol{W} \cup \{I_k\}_{k \in \mathcal{I}})$*, i.e., R cannot cause system variables directly, but only through intervention variables,*

6. $\forall j \in \mathcal{X}: \; R \not\perp_d X_j \implies R \dashrightarrow X_j$*: there are no confounders between R and the system variables (nothing causes R, not even hidden system variables),*

7. $\forall i \in \mathcal{I}, j \in \mathcal{X}: \; I_i \not\perp_d X_j \,|\, R \implies I_i \dashrightarrow X_j$*: there are no confounders between the intervention variables and system variables, when R is in the separating set.*

Proof The propositions follow directly from the JCI assumptions and background knowledge. ■

Adding these background knowledge rules provides a simple means to ruling out several spurious candidate causal structures that do not satisfy the JCI modeling assumptions, showcasing a main advantage of logic-based causal discovery methods. We will refer to the combination of ACID with these rules as ACID-JCI.

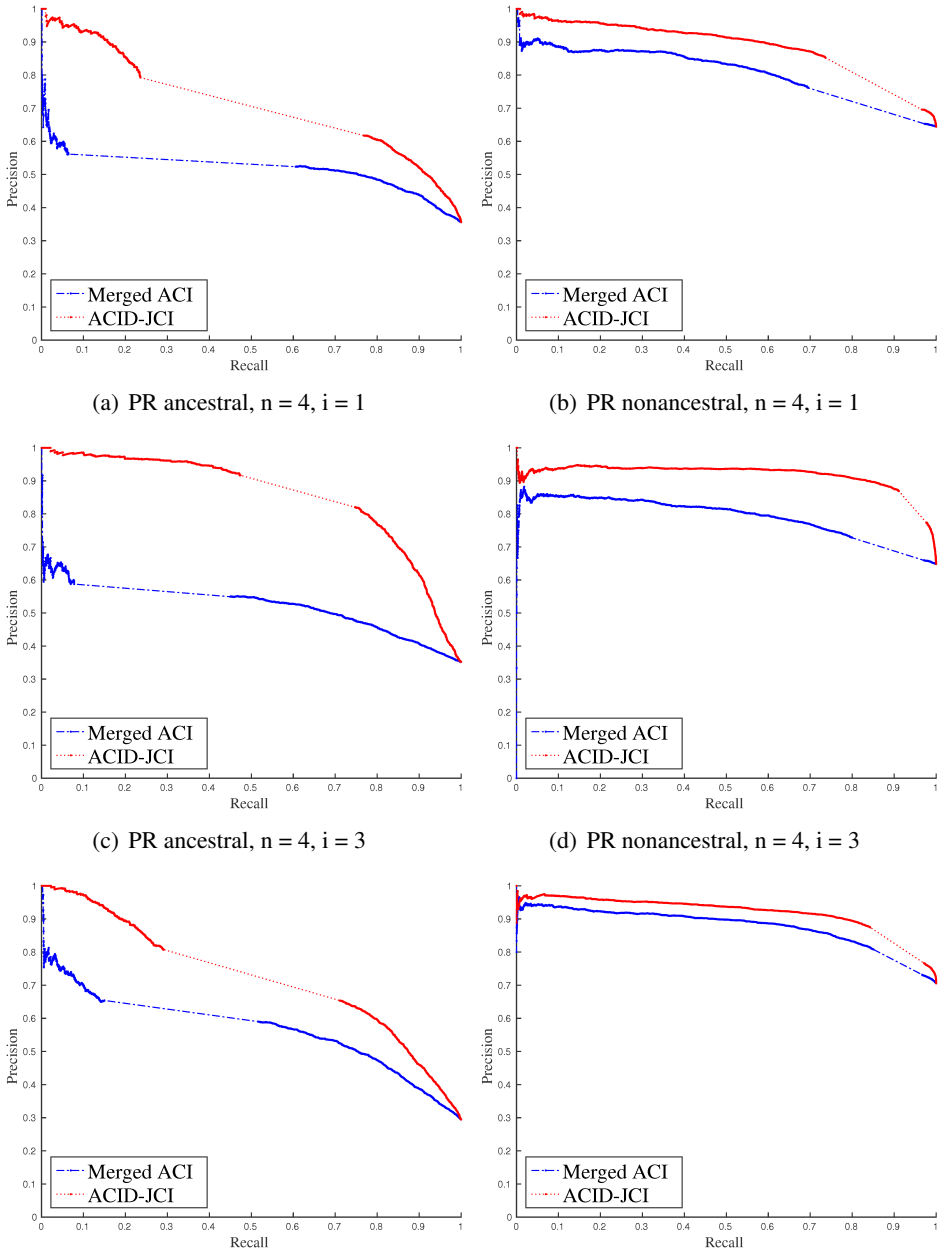

(a) PR ancestral, n = 4, i = 1

(b) PR nonancestral, n = 4, i = 1

(c) PR ancestral, n = 4, i = 3

(d) PR nonancestral, n = 4, i = 3

(e) PR ancestral, n = 6, i = 1

(f) PR nonancestral, n = 6, i = 1

Figure 7: Precision-recall curves on synthetic data for n system variables and i interventions: ancestral predictions (left column) and nonancestral predictions (right column).

7. Evaluation

We evaluate ACID on simulated data in the JCI setting. The simulator builds up on the simulator described in (Hyttinen et al., 2014; Magliacane et al., 2016) and implements soft interventions on unknown targets. For each combination of number of system variables n and interventions i,

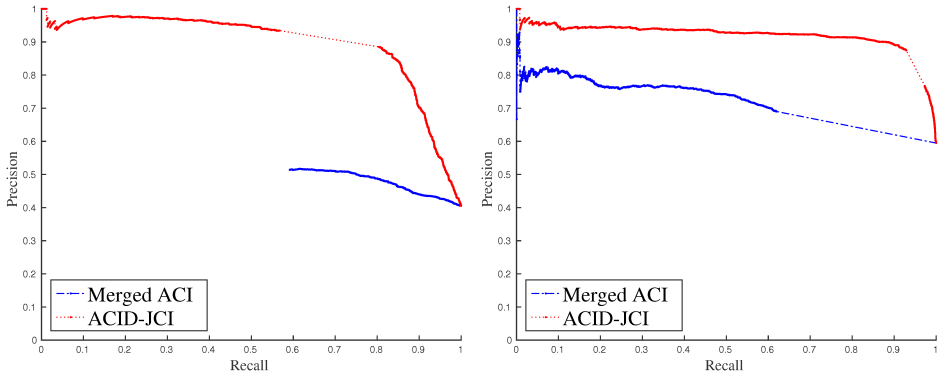

(a) PR ancestral, n = 3, i = 4 (b) PR nonancestral, n = 3, i = 4

Figure 8: Precision-recall curves on synthetic data for $n = 3$ system variables and $i = 4$ interventions: ancestral predictions (left column) and nonancestral predictions (right column).

we generate randomly 1000 linear acyclic models with latent variables and Gaussian noise, and simulate soft interventions on random targets. We then sample $N = 500$ data points for each model, randomly distributed between the i experimental datasets and the observational dataset, perform independence tests and weight the (in)dependence statements using the weighting schemes described in (Magliacane et al., 2016).

In our setting, we evaluate the prediction based on datasets with unknown-target soft interventions. Other existing constraint-based methods like (Hyttinen et al., 2014; Triantafillou and Tsamardinos, 2015) do not apply to this setting as they assume perfect interventions with known targets. Instead, we compare the ancestral structure predicted by ACID-JCI with a naive baseline, in which we merge ancestral structures learnt on each dataset separately with ACI. As inputs to both algorithms we provide the same weighted independence test results, computed with a test based on partial correlation and Fisher's z-transform with significance threshold $\alpha = 0.05$, and the frequentist weighting scheme from (Magliacane et al., 2016).

In Figures 7 and 8 we report the precision and recall (PR) curves for predicting ancestral relations ("indirect" causal relations) and nonancestral relations (the absence of such a causal relation) for different settings of n system variables and i interventions. We can see from the figure that ACID-JCI improves significantly on the accuracy of the predictions with respect to the baseline. As expected, the more interventional datasets available, the better the accuracy for ACID-JCI, as we can see in the case of $n = 4$, $i = 1$ vs. $n = 4$, $i = 3$ in Figure 7. When there are only few datasets, e.g., $i = 1$, the performance of both methods improves as the number of system variables n increases, and the gap between the method gets smaller, as we can see in the case of $n = 4$, $i = 1$ and $n = 6$, $i = 1$ in Figure 7. On the other hand, when there are only few variables, e.g., $n = 3$ in Figure 8, ACID-JCI is able to predict correctly several ancestral relations that are not predicted otherwise, illustrating that the *Joint* Causal Inference framework not only leads to statistical advantages but also enlarges the set of identifiable ancestral relations compared to methods that deal with each dataset separately.

8. Conclusions and future work

In this paper, we presented Joint Causal Inference (JCI), a powerful formulation of causal discovery over multiple datasets that was previously unexploited by constraint-based methods. Current constraint-based methods cannot be applied out-of-the-box to JCI because of faitfulness violations, so we proposed a simple strategy for dealing with this type of faithfulness violations. We implement this strategy in ACID, a determinism-tolerant extension of a recently proposed causal discovery method, and apply ACID to JCI, showing its benefits in an evaluation on simulated data.

In JCI the assumption of a unique underlying causal DAG precludes certain types of interventions. There are several techniques to extend our formulation of the problem to perfect interventions, or other interventions that induce new independences. For example, given an observational dataset, we could identify the datasets with perfect interventions by noticing the additional independences, perform causal inference on each of them separately and merge the predictions in a similar way to (Hyttinen et al., 2014; Triantafillou and Tsamardinos, 2015). In future work, we plan to investigate these techniques, as well as techniques to include fat-hand interventions that induce new dependences between the intervention targets.

Moreover, we plan to investigate other possible strategies or extensions to existing algorithms for dealing with faithfulness violations due to deterministic relations. Finally, although very accurate and flexible, logic-based methods as (Hyttinen et al., 2014; Magliacane et al., 2016) are limited in the number of possible variables they can handle. JCI introduces additional variables, reducing their scalability even more. We plan to investigate improvements to the execution times of methods like ACID.

ACKNOWLEDGMENTS

SM, JMM and TC were supported by NWO, the Netherlands Organization for Scientific Research (VIDI grant 639.072.410). SM was also supported by the Dutch programme COMMIT/ under the Data2Semantics project. TC was also supported by NWO grant 612.001.202 (MoCoCaDi), and EU-FP7 grant agreement n.603016 (MATRICS).

REFERENCES

Giorgos Borboudakis and Ioannis Tsamardinos. Towards robust and versatile causal discovery for business applications. *KDD*, 2016.

Diego Colombo, Marloes H. Maathuis, Markus Kalisch, and Thomas S. Richardson. Learning high-dimensional directed acyclic graphs with latent and selection variables. *The Annals of Statistics*, 40(1): 294–321, 2012.

P. Constantinou and A. P. Dawid. Extended Conditional Independence and Applications in Causal Inference. *ArXiv e-prints*, December 2015.

Gregory F. Cooper. A simple constraint-based algorithm for efficiently mining observational databases for causal relationships. *Data Min. Knowl. Discov.*, 1(2):203–224, January 1997. ISSN 1384-5810. doi: 10.1023/A:1009787925236. URL http://dx.doi.org/10.1023/A:1009787925236.

A. P. Dawid. Influence diagrams for causal modelling and inference. *International Statistical Review*, 70 (2):161–189, 2002. ISSN 1751-5823. doi: 10.1111/j.1751-5823.2002.tb00354.x. URL http://dx.doi.org/10.1111/j.1751-5823.2002.tb00354.x.

A. Philip Dawid. Conditional independence in statistical theory. *Journal of the Royal Statistical Society B*, 41(1):1–31, 1979.

Steffen Ducheyne. Galileo's interventionist notion of "cause". *Journal of the History of Ideas*, 67(3): 443–464, 2006. ISSN 00225037, 10863222. URL http://www.jstor.org/stable/30141036.

D. Eaton and K. Murphy. Exact Bayesian structure learning from uncertain interventions. In *Proceedings of the 10th International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 107–114, 2007.

Frederick Eberhardt. A sufficient condition for pooling data. *Synthese*, 163(3):433–442, 2008. ISSN 1573-0964. doi: 10.1007/s11229-007-9293-3. URL http://dx.doi.org/10.1007/s11229-007-9293-3.

Doris Entner, Patrik O. Hoyer, and Peter Spirtes. Data-driven covariate selection for nonparametric estimation of causal effects. In *AISTATS*, volume 31 of *JMLR Proceedings*, pages 256–264. JMLR.org, 2013.

Dan Geiger, Thomas Verma, and Judea Pearl. Identifying independence in Bayesian networks. *Networks*, 20(5):507–534, 1990.

Alain Hauser and Peter Bühlmann. Characterization and greedy learning of interventional Markov equivalence classes of directed acyclic graphs. *Journal of Machine Learning Research*, 13(Aug):2409–2464, 2012.

Antti Hyttinen, Frederick Eberhardt, and Matti Järvisalo. Constraint-based causal discovery: Conflict resolution with answer set programming. In *UAI 2014*, pages 340–349, 2014.

Patrick Kemmeren, Katrin Sameith, Loes A.L. van de Pasch, Joris J. Benschop, Tineke L. Lenstra, Thanasis Margaritis, Eoghan O'Duibhir, Eva Apweiler, Sake van Wageningen, Cheuk W. Ko, Sebastiaan van Heesch, Mehdi M. Kashani, Giannis Ampatziadis-Michailidis, Mariel O. Brok, Nathalie A.C.H. Brabers, Anthony J. Miles, Diane Bouwmeester, Sander R. van Hooff, Harm van Bakel, Erik Sluiters, Linda V. Bakker, Berend Snel, Philip Lijnzaad, Dik van Leenen, Marian J.A. Groot Koerkamp, and Frank C.P. Holstege. Large-scale genetic perturbations reveal regulatory networks and an abundance of gene-specific repressors. *Cell*, 157(3):740 – 752, 2014. ISSN 0092-8674. doi: http://dx.doi.org/10.1016/j.cell.2014.02.054. URL http://www.sciencedirect.com/science/article/pii/S0092867414003420.

Vincenzo Lagani, Ioannis Tsamardinos, and Sofia Triantafillou. *Artificial Intelligence: Theories and Applications: 7th Hellenic Conference on AI, SETN 2012, Lamia, Greece, May 28-31, 2012. Proceedings*, chapter Learning from Mixture of Experimental Data: A Constraint–Based Approach. 2012.

Jan Lemeire, Stijn Meganck, Francesco Cartella, and Tingting Liu. Conservative independence-based causal structure learning in absence of adjacency faithfulness. *Int. J. Approx. Reasoning*, 53(9):1305–1325, December 2012. ISSN 0888-613X. doi: 10.1016/j.ijar.2012.06.004. URL http://dx.doi.org/10.1016/j.ijar.2012.06.004.

Sara Magliacane, Tom Claassen, and Joris M. Mooij. Ancestral causal inference. In *NIPS*, 2016.

F. Markowetz, S. Grossmann, and R. Spang. Probabilistic soft interventions in conditional Gaussian networks. In *Tenth International Workshop on Artificial Intelligence and Statistics (AISTATS)*, pages 214–221, 2005.

Joris M. Mooij and Tom Heskes. Cyclic causal discovery from continuous equilibrium data. In Ann Nicholson and Padhraic Smyth, editors, *Proceedings of the 29th Annual Conference on Uncertainty in Artificial Intelligence (UAI-13)*, pages 431–439. AUAI Press, 2013. URL http://auai.org/uai2013/prints/papers/23.pdf.

J. Pearl. *Causality: models, reasoning and inference*. Cambridge University Press, 2009.

J. Peters, P. Bühlmann, and N. Meinshausen. Causal inference using invariant prediction: identification and confidence intervals. *Journal of the Royal Statistical Society, Series B*, 2015.

Jonas Peters, Joris M. Mooij, Dominik Janzing, Bernhard Schölkopf, et al. Causal discovery with continuous additive noise models. *Journal of Machine Learning Research*, 15(1):2009–2053, 2014.

P. Spirtes, C. Glymour, and R. Scheines. *Causation, Prediction, and Search*. MIT press, 2nd edition, 2000.

J. Tian and J. Pearl. Causal discovery from changes. In *Proceedings of the Seventeenth Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 512–521, 2001.

Sofia Triantafillou and Ioannis Tsamardinos. Constraint-based causal discovery from multiple interventions over overlapping variable sets. *Journal of Machine Learning Research*, 16:2147–2205, 2015.

Jiji Zhang. On the completeness of orientation rules for causal discovery in the presence of latent confounders and selection bias. *Artif. Intell.*, 172(16-17):1873–1896, November 2008. ISSN 0004-3702. doi: 10.1016/j.artint.2008.08.001. URL `http://dx.doi.org/10.1016/j.artint.2008.08.001`.

9. Proofs

For completeness we provide the proofs for the ACID rules, which are just slight modifications of the proofs for the ACI (Magliacane et al., 2016) rules. These proofs were already implicitly based on d-separation concepts, with two additional implicit steps that were based on the Causal Markov and Faithfulness Assumptions: one in which each conditional independence implies a d-separation, and another one in which each d-separation implies a conditional independence. In this reformulation we can make explicit these steps and decouple them from the rules.

Lemma 14 *For X, Y, Z, U, $\boldsymbol{W}$ disjoint (sets of) variables of a DAG (or ADMG) $\mathcal{G}$:*

1. $(X \perp_d Y \mid \boldsymbol{W}) \wedge (X \not\dashrightarrow \boldsymbol{W}) \implies X \not\dashrightarrow Y$
2. $(X \perp_d Y \mid \boldsymbol{W}) \wedge (X \not\perp_d Y \mid \boldsymbol{W} \cup Z) \implies (X \not\perp_d Z \mid \boldsymbol{W}) \wedge (Z \not\dashrightarrow \{X, Y\} \cup \boldsymbol{W})$
3. $(X \not\perp_d Y \mid \boldsymbol{W}) \wedge (X \perp_d Y \mid \boldsymbol{W} \cup Z) \implies (X \not\perp_d Z \mid \boldsymbol{W}) \wedge (Z \dashrightarrow \{X, Y\} \cup \boldsymbol{W})$
4. $(X \not\perp_d Y \mid \boldsymbol{W}) \wedge (X \perp_d Y \mid \boldsymbol{W} \cup Z) \wedge (X \perp_d Z \mid \boldsymbol{W} \cup U) \implies X \perp_d Y \mid \boldsymbol{W} \cup U$
5. $(Z \not\perp_d X \mid \boldsymbol{W}) \wedge (Z \not\perp_d Y \mid \boldsymbol{W}) \wedge (X \perp_d Y \mid \boldsymbol{W}) \implies X \not\perp_d Y \mid \boldsymbol{W} \cup Z$

Proof

1. This is a strengthened version of rule $\mathcal{R}2$(i) in (Entner et al., 2013): note that the additional assumptions made there ($Y \not\dashrightarrow \boldsymbol{W}$, $Y \not\dashrightarrow X$) are redundant and not actually used in their proof. For completeness, we give the proof here. If $X \dashrightarrow Y$, then there is a directed path from X to Y in $\mathcal{G}$. As all paths between X and Y in $\mathcal{G}$ are blocked by $\mathbf{W}$, the directed path from X to Y must contain a node $W \in \boldsymbol{W}$. Hence $X \dashrightarrow W$, a contradiction with $X \not\dashrightarrow \mathbf{W}$.

2. If $(X \perp_d Y \mid \boldsymbol{W}) \wedge (X \not\perp_d Y \mid \boldsymbol{W} \cup Z)$ then there exists a path π in $\mathcal{G}$ between X and Y such that each noncollider on π is not in $\boldsymbol{W} \cup \{Z\}$, every collider on π is ancestor of $\boldsymbol{W} \cup \{Z\}$, and there exists a collider on π that is ancestor of Z but not of $\boldsymbol{W}$. Let C be the collider on π closest to X that is ancestor of Z but not of $\boldsymbol{W}$. Note that

   (a) The path $X \cdots C \to \cdots \to Z$ is d-connected given $\boldsymbol{W}$.

   (b) $Z \not\dashrightarrow \boldsymbol{W}$ (because otherwise $C \dashrightarrow Z \dashrightarrow \boldsymbol{W}$, a contradiction).

(c) $Z \not\dashrightarrow Y$ (because otherwise the path $X \cdots C \to \cdots \to Z \to \cdots \to Y$ would be d-connected given $\boldsymbol{W}$, a contradiction).

Hence we conclude that $X \not\perp_d Z \mid \boldsymbol{W}$, $Z \not\dashrightarrow \boldsymbol{W}$, $Z \not\dashrightarrow Y$, and by symmetry also $Z \not\dashrightarrow X$.

3. Suppose $(X \not\perp_d Y \mid \boldsymbol{W}) \wedge (X \perp_d Y \mid \boldsymbol{W} \cup Z)$. Then there exists a path π in $\mathcal{G}$ between X and Y, such that each noncollider on π is not in $\boldsymbol{W}$, each collider on π is an ancestor of $\boldsymbol{W}$, and Z is a noncollider on π. Note that

   (a) The subpath $X \dots Z$ must be d-connected given $\boldsymbol{W}$.

   (b) Z has at least one outgoing edge on π. Follow this edge further along π until reaching either X, Y, or the first collider. When a collider is reached, follow the directed path to $\boldsymbol{W}$. Hence there is a directed path from Z to X or Y or to $\boldsymbol{W}$, i.e., $Z \dashrightarrow \{X, Y\} \cup \boldsymbol{W}$.

4. If in addition, $X \perp_d Z \mid \boldsymbol{W} \cup U$, then U must be a noncollider on the subpath $X \dots Z$. Therefore, $X \perp_d Y \mid \boldsymbol{W} \cup U$.

5. Assume that $Z \not\perp_d X \mid \boldsymbol{W}$ and $Z \not\perp_d Y \mid \boldsymbol{W}$. Then there must be paths π between Z and X and ρ between Z and Y in $\mathcal{G}$ such that each noncollider is not in $\boldsymbol{W}$ and each collider is ancestor of $\boldsymbol{W}$. Let U be the node on π closest to X that is also on ρ (this could be Z). Then we have a path $X \cdots U \cdots Y$ such that each collider (except U) is ancestor of $\boldsymbol{W}$ and each noncollider (except U) is not in $\boldsymbol{W}$. This path must be blocked given $\boldsymbol{W}$ as $X \perp_d Y \mid \boldsymbol{W}$. If U would be a noncollider on this path, it would need to be in $\boldsymbol{W}$ in order to block it; however, it must then also be a noncollider on π or ρ and hence cannot be in $\boldsymbol{W}$. Therefore, U must be a collider on this path and cannot be ancestor of $\boldsymbol{W}$. We have to show that U is ancestor of Z. If U were a collider on π or ρ, it would be ancestor of $\boldsymbol{W}$, a contradiction. Hence U must have an outgoing arrow pointing towards Z on π and ρ. If we encounter a collider following the directed edges, we get a contradiction, as that collider, and hence U, would be ancestor of $\boldsymbol{W}$. Hence U is ancestor of Z, and therefore, $X \not\perp_d Y \mid \boldsymbol{W} \cup Z$.

■